EC'18 Tutorial: The Menu-Size Complexity of Precise and Approximate Revenue-Maximizing Auctions

Kira Goldner[1] and Yannai A. Gonczarowski[2]

[1]University of Washington
[2]The Hebrew University of Jerusalem and Microsoft Research

Cornell University, Ithaca, NY, June 18, 2018

Schedule:

08:30am – 10:30am: The Menu-Size of Multi-Item Auctions (Yannai)
11:00am – 12:30am: The Menu-Size of FedEx and Related Auctions (Kira)

The Menu-Size of Multi-Item Auctions

Yannai A. Gonczarowski

The Hebrew University of Jerusalem and Microsoft Research

EC'18 Tutorial: The Menu-Size of Precise and Approximate Revenue-Maximizing Auctions

Cornell University, Ithaca, NY, June 18, 2018

This Lecture

- Aiming to not assume any prior knowledge of menu-size concepts.
- Shout out if you have a question.
- Tutorial goal: start from basics, and get you acquainted with the recent explosion of results, directions, and open questions on menu sizes.
  - This lecture: multi-item auctions.
  - Kira's lecture (after the break): FedEx and related auctions.
- Focus on results, with only glimpses of proofs/techniques (mostly for intuition regarding open questions).
- Lecture/tutorial order not chronological.
- Did I miss a relevant result? Please talk to me / email me.

Model

- A single **seller** has n nonidentical **items** that she would like to sell. The seller has no other use (and has no cost) for the items.
- There is one (potential) **buyer**, who has a private **value** (maximum willingness to pay) for each item (need not be the same for all items).
- The buyer's valuation is **additive**: her value for any subset of the items is the **sum** of her values for the items in the subset.
- The buyer's utility is **quasilinear**: her total utility is the sum of her values for the items that she holds, minus any payments she has made.
- The buyer has no budget constraints.
- Stylized model: the seller knows a **prior distribution** over the buyer's values for the individual items. (The values for the various item may be correlated.)
- (The buyer and seller are **risk-neutral**: seller cares only about expected revenue, buyer cares only about expected utility.)

Auctions / Mechanisms

- A (direct-revelation) auction mechanism is a function that maps (in a possibly randomized manner) each possible buyer **type** (specification of value for each item) to an **outcome**: which items to award the buyer, and how much to charge the buyer.
- A mechanism is **individually rational (IR)** if the buyer can always opt out: the expected utility (over the randomness of the mechanism) of any buyer type is always nonnegative. (Mechanism cannot be "pay me a billion dollars and get Item 3.")
- A mechanism is **incentive compatible (IC)** if the buyer has no incentive to **strategize**: $\forall t, t' : \mathbb{E}\left[u_t(M(t))\right] \geq \mathbb{E}\left[u_t(M(t'))\right]$, where the expectation is over the randomness of the mechanism. (Mechanism cannot be "tell me your type, now take all items and pay me your value for all items.")
- The seller wishes to choose a **truthful** (IR+IC) mechanism that maximizes her expected revenue, where the expectation is over both the prior distribution and the randomness of the mechanism.

One Item

- A possible mechanism: choose a price, and offer the item for that price.
- Among all posted-price mechanisms, the one obtaining highest revenue is the one posting a price of

$$\arg\mathrm{Max}_p\, p \cdot \mathbb{P}_{v\sim F}\left[v \geq p\right].$$

- Other mechanisms also possible (e.g., lottery tickets).

Theorem (Myerson, 1981; Riley and Zeckhauser, 1983)

No other mechanism can obtain better revenue than posting the revenue-maximizing price.

Two Items

How can the seller maximize the revenue from two items?

- If independent, optimally sell each item separately? ✗

Example

If both item values are uniformly distributed in $\{1, 2\}$:

- Pricing each item separately, seller obtains a revenue of \$1 for each item, for a total revenue of \$2.
- Pricing only the bundle at \$3, seller obtains a revenue of $\$3 \cdot 0.75 = 2.25 \ > \ 2$!
- So pricing each item separately does not always maximize revenue!

Two Items

How can the seller maximize the revenue from two items?

- If independent, optimally sell each item separately? ✗
- Optimally sell the bundle of both items? ✗
- Either sell separately or bundle? ✗
- Post a price for each item and a price for the bundle? ✗
- Choose between a few lotteries? ✗

Distribution	Optimal Mechanism
$U(\{1,2\}) \times U(\{1,2\})$	Sell the bundle (for \$3)
$U(\{0,1\}) \times U(\{0,1\})$	Sell each separately (\$1 each)
$U(\{0,1,2\}) \times U(\{0,1,2\})$	Offer: one for \$2 / both for \$3
$U(\{1,2\}) \times U(\{1,3\})$	Offers include lottery tickets (both for \$4 / for \$2.5: first w.p. 1, second w.p. 1/2) T'04,DDT'14
$\mathrm{Beta}(1,2) \times \mathrm{Beta}(1,2)$	Offer infinitely many lotteries DDT'13,DDT'15

Not Merely Unaesthetic / Hard to Formally Analyze

- Hard (#P-Hard) to compute. DDT'14
- Harder to represent to the participant.
- Harder for the participant to find/verify optimal strategy.

So what can we get using simpler auctions?

Simple Auctions: Limiting Complexity

Option 1: Qualitatively: disallow some "features":

- Allow only separate selling. HN'12
- Allow only "packaging". BILW'14, R'16
- Disallow lotteries. BNR'18

An "all or nothing" approach...

Such studied features lose at least a constant factor of revenue.

Option 2: Quantitatively: limit a numeric complexity measure:

- Number of options presented to the buyer. HN'13
- Length of auction description using any language. DHN'14
- Learning-theoretic dimensionality. MR'15, MR'16, BSV'16, BSV'18

A " " approach...

This tutorial.

The Menu Size of an Auction Mechanism

- By the **Taxation Principle**, every truthful mechanism, however complex, is equivalent to specifying a **menu** of possible probabilistic outcomes for the buyer to choose from.

Chez Seller

Items • Bundles • Lotteries

Today's Specials

$\mathbb{P}$[Item 1]	$\mathbb{P}$[Item 2]	$\mathbb{E}$[Price]
0%	100%	\$3
20%	30%	\$4
40%	60%	\$10
⋮	⋮	⋮
100%	100%	\$20

Menu Size
Hart-Nisan'13

The Classic Choice

0%	0%	\$0

— One entry per buyer —

- Was floating around as a proof technique even before 2013: Briest-Chawla-Kleinberg-Weinberg'10, Dobzinski'11, Dughmi-Vondrak'11, Dobzinski-Vondrak'12.

Menu-Size as Complexity Measure: Pros and Cons

Pros:

- Simple and intuitive to define.
- Tractable to analyze.
- The base-2 logarithm (rounded up) of the menu-size is the deterministic **communication complexity** of computing the auction outcome (Babaioff-G.-Nisan'17).

Cons:

- There may be auctions that are intuitively simple, and also concise to describe, but have a large menu size.
- Main example: selling n items separately; menu-size $\exp(n)$.
  - Indeed, high (linear in n) communication complexity… but still seems "simple." Mitigating: separate-selling revenue attainable via poly(n) menu size (Babaioff-G.-Nisan'17).
  - Switch to "additive menu" size? (Definition w.r.t. lotteries?)
  - By any natural definition, even for two i.i.d. items, the optimal revenue cannot be attained by an "additive menu" (Babaioff-Nisan-Rubinstein'18).

Auction complexity measures trade-off simplicity of definition (e.g., menu size) with flexibility (e.g., Kolmogorov complexity).

Menu-Size Complexity

- The menu size is extremely simple and intuitive to define, and directly implies the communication complexity of running the mechanism.
- But it is important to be aware of its flaws (esp. as a function of the number of item n; actually as a function of other parameters such as ε or H that will be defined later, the above criticism does not directly apply, or even does not apply at all).
- Important to analyze this simple measure; important to understand other auction complexity measures as well. Must start from somewhere...
- Maybe someone in the audience will suggest a new measure for auction complexity?

Precise Revenue Maximization

Theorem (Daskalakis et al., 2013, 2015)

There exists a distribution $F \in \Delta([0,1])$ s.t. the menu of the optimal mechanism for $F \times F$ has a continuum of menu entries.

- $F = \text{Beta}(1,2)$: distributed over $[0,1]$ w/density $2(1-x)$.
- Proof uses their optimal-transport duality framework.

So, for precise revenue maximization:

- One item: menu-size 1 suffices (Myerson'81, RZ'83).
- Two items, even bounded, i.i.d., "nice" distributions: infinite menu-size required.

Two ways to proceed from here:

- Approximate revenue maximization — rest of this lecture.
- Find a model "in between" one item and two i.i.d. items. — Kira's lecture after the break.

Approximate Revenue Maximization

Theorem (Hart-Nisan'13, inspired by proof of Briest-Chawla-Kleinberg-Weinberg'10 for $n \geq 3$)

For every number of items $n \geq 2$, every $\varepsilon > 0$, and every menu-size m, there exists a distribution $F \in \Delta\big([0,1]^n\big)$ s.t.

$$\underset{\text{an auction with menu-size } m}{\text{Optimal revenue attainable by}} \Rightarrow \mathcal{R}ev_{[m]}(F) < \varepsilon \cdot \mathcal{R}ev(F) \Leftarrow \underset{\text{by any (truthful) auction}}{\text{Optimal revenue attainable}}$$

- In particular, deterministic mechanisms cannot guarantee any fraction of the optimal revenue.
- Compare: Hart-Reny'17 (see also Hart-Nisan'12): selling two independent items separately attains $\geq 62\%$ of OPT.

Three ways to proceed from here:

- How does the revenue improve with the menu size?
- Relax our goal: additive approximation.
- Restrict distribs.: bounded also from below / independent.

Will touch on all above, focus on independent item distributions.

Revenue Improvement with Menu Size

Theorem (Hart-Nisan'13)

For every $n \geq 2$ and every $F \in \Delta(\mathbb{R}^n_+)$:

1. $\mathcal{R}ev_{[1]}(F) = \mathcal{B}Rev(F) \Leftarrow$ *Revenue attainable by optimally pricing the bundle of all items*
2. $\mathcal{R}ev_{[m_1+m_2]}(F) \leq \mathcal{R}ev_{[m_1]}(F) + \mathcal{R}ev_{[m_2]}(F)$ *for all m_1, m_2.*
3. $\mathcal{R}ev_{[m]}(F) \leq m \cdot \mathcal{R}ev_{[1]}(F)$ *for all m.*

How tight is Part 3? Obviously, for some F bundling is optimal, so for such distributions $\mathcal{R}ev_{[m]}(F) = \mathcal{R}ev_{[1]}(F)$ for all m.

Theorem (Hart-Nisan'13)

For every $n \geq 2$, there exists a distribution $F \in \Delta(\mathbb{R}^n_+)$ with $\mathcal{R}ev_{[1]}(F) \in (0, \infty)$ s.t.

$$\mathcal{R}ev_{[m]}(F) \geq \Omega(m^{1/7}) \cdot \mathcal{R}ev_{[1]}(F).$$

They conjecture that the constant $1/7$ can be improved upon...

Additive Approximation for Bounded Domains

Theorem (Dughmi-Han-Nisan'14, see also Hart-Nisan'13)

There exists $C(n,\varepsilon) = \left({}^{\log n}/_{\varepsilon}\right)^{O(n)}$ s.t. for every n, every $\varepsilon > 0$, and every $F \in \Delta\big([0,1]^n\big)$,

$$\mathcal{R}ev_{[C(n,\varepsilon)]}(F) > \mathcal{R}ev(F) - \varepsilon$$

- An upper bound! Proof technique: **"nudge and round."**
- We will prove that $C(n,\varepsilon) = \left({}^{n}/_{\varepsilon}\right)^{O(n)}$ (Hart-Nisan'13):
  1. Start with the optimal menu.
  2. **Nudge:** discount all prices multiplicatively: $p \leftarrow (1 - {}^{\varepsilon}/_{n}) \cdot p$.
  3. Discretize by **rounding** probabilities (could have as well rounded price) to multiples of ${}^{\varepsilon^2}/_{n^2}$. (DHN'14: log grid.)
- Revenue loss at most 2ε. Indeed, if the original payment by some buyer type is p, then new payment $\geq (1 - {}^{\varepsilon}/_{n})(p - \varepsilon)$:
  - Post-discounting, the utility from any menu-entry originally costing less than $p - \varepsilon$ is at least ${}^{\varepsilon^2}/_{n}$ less than the utility from the originally chosen menu-entry.
  - Discretizing decreased the utility from the originally chosen entry by at most ${}^{\varepsilon^2}/_{n}$, so could not have tilted the balance.

Multiplicative Loss vs. Additive Loss

Recall that on one hand:

Theorem (Hart-Nisan'13, see also Briest-Chawla-Kleinberg-Weinberg'10)

For every number of items $n \geq 2$, every $\varepsilon > 0$, and every menu-size m, there exists a distribution $F \in \Delta\big([0,1]^n]\big)$ s.t.

$$\mathcal{R}ev_{[m]}(F) < \varepsilon \cdot \mathcal{R}ev(F)$$

And on the other hand:

Theorem (Dughmi-Han-Nisan'14, see also Hart-Nisan'13)

There exists $C(n,\varepsilon) = \big({}^{\log n}/_{\varepsilon}\big)^{O(n)}$ s.t. for every n, every $\varepsilon > 0$, and every $F \in \Delta\big([0,1]^n\big)$,

$$\mathcal{R}ev_{[C(n,\varepsilon)]}(F) > \mathcal{R}ev(F) - \varepsilon$$

So the impossibility in the first theorem above comes from the case of very small optimal revenues.

Multiplicative Loss with Bounded Support

Indeed, the above additive upper bounds follow from:

Theorem (Dughmi-Han-Nisan'14, see also Hart-Nisan'13)

There exists $C(n,\varepsilon,H) = \left(\frac{\log n + \log H}{\varepsilon}\right)^{O(n)}$ s.t. for every n, every $\varepsilon > 0$, every H and every $F \in \Delta\big([1,H]^n\big)$,

$$\mathcal{R}ev_{[C(n,\varepsilon,H)]}(F) > (1-\varepsilon)\cdot\mathcal{R}ev(F)$$

Minimum of 1 is w.l.o.g., since can scale $[L,H]$ to $[1, {}^{H}/_{L}]$.

Theorem (Dughmi-Han-Nisan'14)

For distributions supported on $[1,H]^n$:

1. *Menu-size n can attain a $\Omega({}^{1}/_{\log H})$ fraction of the revenue.*
2. *Auctions with **Kolmogorov complexity** polynomial in n guarantee at most a $O({}^{1}/_{\log H})$ fraction of the revenue.*

$\Rightarrow$ Moving to any fancier complexity measure will not help here.

Guarantees for Unbounded Distributions

For unbounded distribs., only multiplicative loss makes sense.
On one hand:

Theorem (Hart-Nisan'13, see also Briest et al.'10)

For every $n \geq 2$, every $\varepsilon > 0$, and every menu-size m, there exists a distribution $F \in \Delta\big([0,1]^n\big)$ s.t.

$$\mathcal{R}ev_{[m]}(F) < \varepsilon \cdot \mathcal{R}ev(F)$$

And on the other hand:

Theorem

For every n, every $\varepsilon > 0$, and every L, H, there exists $C(n, \varepsilon, {}^{L}\!/_{H}) = \big(\frac{\log n + \log {}^{L}\!/_{H}}{\varepsilon}\big)^{O(n)}$ s.t. for every $F \in \Delta\big([L,H]^n\big)$,

$$\mathcal{R}ev_{[C(n,\varepsilon,{}^{L}\!/_{H})]}(F) > (1-\varepsilon) \cdot \mathcal{R}ev(F)$$

What about restricting the distributions in some way other than bounding? (In any way, "nudge and round" no longer suffices.)

Independent Values

Theorem (Babaioff-G.-Nisan'17)

For every n and every $\varepsilon > 0$, there exists $C(n,\varepsilon) = \left(\log n/\varepsilon\right)^{O(n)}$ s.t. for every $F_1,\ldots,F_n \in \Delta(\mathbb{R}_+)$,

$$\mathcal{R}ev_{[C(n,\varepsilon)]}(F_1 \times \cdots \times F_n) > (1-\varepsilon)\cdot\mathcal{R}ev(F_1 \times \cdots \times F_n).$$

- Recall: "nudge & round" not suitable for unbounded distributions. (Grid discrete but infinite.)
- Rough high-level overview:
  - Scale so that $\mathcal{R}ev(F) = 1$.
  - For suitable $H = \text{poly}(n,\varepsilon)$:
- Main thing to note: exponential menu size only due to selling to the core.

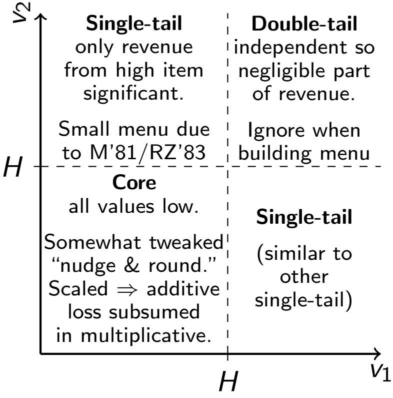

Required Menu Size

Theorem (Babaioff-G.-Nisan'17)

For every n and every $\varepsilon > 0$, there exists $C(n,\varepsilon) = \left(\log n/\varepsilon\right)^{O(n)}$ s.t. for every $F_1,\ldots,F_n \in \Delta(\mathbb{R}_+)$,

$$\mathcal{R}ev_{[C(n,\varepsilon)]}(F_1 \times \cdots \times F_n) > (1-\varepsilon)\cdot\mathcal{R}ev(F_1 \times \cdots \times F_n).$$

How fast must C grow as a function of n and ε?

Theorem (Babaioff-Immorlica-Lucier-Weinberg'14)

For product distributions, either selling separately or selling bundled guarantees at least $c > 1/6$ of the optimal revenue.

Theorem (Babaioff-G.-Nisan'17)

For product distributions, the revenue from selling separately can be attained up to a multiplicative ε via menu-size $n^{d(\varepsilon)}$.

$\Rightarrow$ poly(n) menu-size guarantees $1/6$ of optimal revenue.

Required Menu Size

Theorem (Babaioff-G.-Nisan'17)

For every n and every $\varepsilon > 0$, there exists $C(n,\varepsilon) = \left({}^{\log n}/_{\varepsilon}\right)^{O(n)}$ s.t. for every $F_1, \ldots, F_n \in \Delta(\mathbb{R}_+)$,

$$\mathcal{R}ev_{[C(n,\varepsilon)]}(F_1 \times \cdots \times F_n) > (1-\varepsilon) \cdot \mathcal{R}ev(F_1 \times \cdots \times F_n).$$

How fast must C grow as a function of n and ε?

Theorem (Babaioff-G.-Nisan'17, see also Babaioff et al.'14)

There exists d s.t. for every n and every $F_1, \ldots, F_n \in \Delta(\mathbb{R}_+)$,

$$\mathcal{R}ev_{[n^d]}(F_1 \times \cdots \times F_n) > {}^1/_6 \cdot \mathcal{R}ev(F_1 \times \cdots \times F_n).$$

Theorem (Babaioff-G.-Nisan'17)

Fix $F = U\big(\{0,1\}\big)^n$, then $\mathcal{R}ev_{[2^{n/10}]}(F) < (1 - \frac{1}{10n}) \cdot \mathcal{R}ev(F)$.

Note: can sell the bundle w.h.p. and lose $\approx {}^1/\sqrt{n}$ of the revenue.

Independent: Revenue Guarantee vs. Menu Size

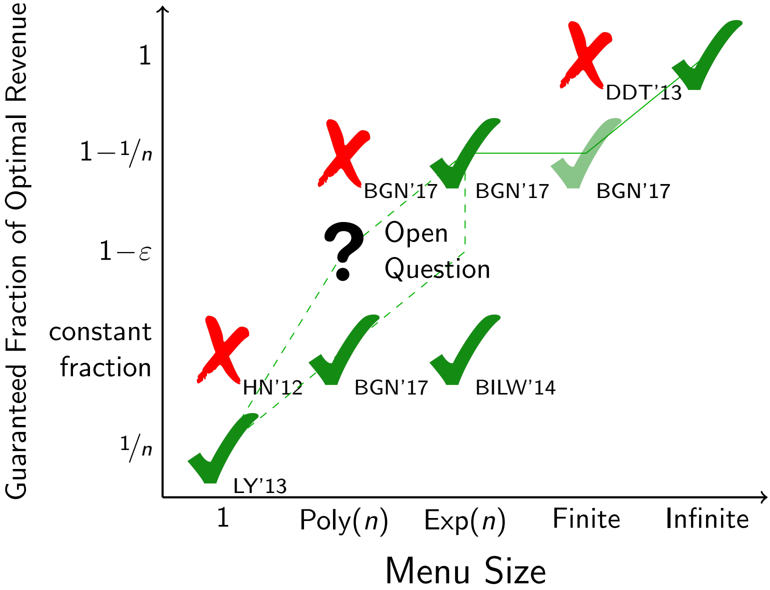

Correlated: Revenue Guarantee vs. Menu Size

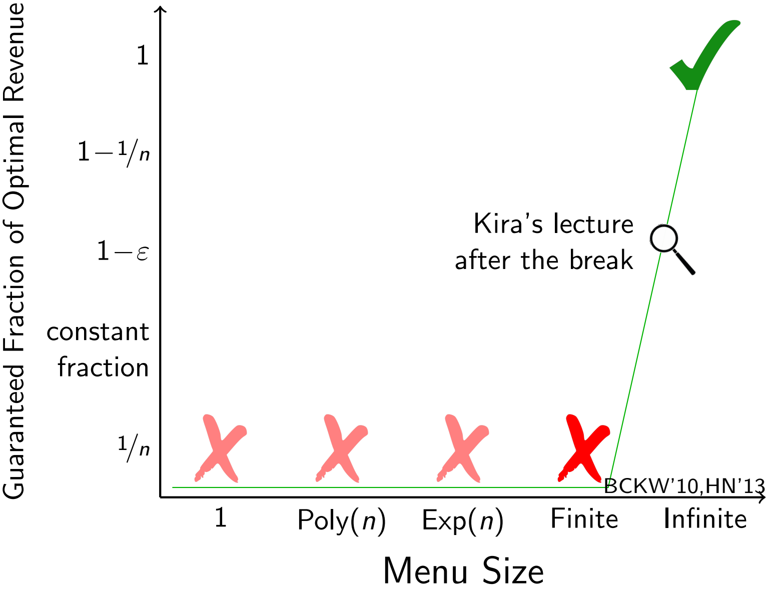

Dependence on ε: Lower Bound for Two Items

DDT'13,15 $\Rightarrow$ required menu-size is $\omega(1)$ as a function of ε.

Theorem (G.'18)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ *and* $F \in \Delta\big([0,1]\big)$, *s.t. for every* $\varepsilon > 0$,

$$\mathcal{R}ev_{[C(\varepsilon)]}(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$$

- $F = \text{Beta}(1,2)$ as in DDT. (Also via optimal transport.)
- $\Rightarrow$ same lower bound for **multiplicative** ε loss, even for i.i.d.
- $\Rightarrow$ same $\Omega(1/\sqrt[4]{\varepsilon})$ lower bound for any fixed n.
- $\Rightarrow$ For fixed n, menu-size $\text{poly}(1/\varepsilon)$ necessary and sufficient.

Dependence on ε: Communication Complexity

DDT'13,15 $\Rightarrow$ required menu-size is $\omega(1)$ as a function of ε.

Theorem (G.'18)

There exist $C(\varepsilon) = \Omega(1/\sqrt[4]{\varepsilon})$ and $F \in \Delta\big([0,1]\big)$, s.t. for every $\varepsilon > 0$,

$$\mathcal{R}ev_{[C(\varepsilon)]}(F \times F) < \mathcal{R}ev(F \times F) - \varepsilon$$

- $\Rightarrow$ For fixed n, menu-size poly$(1/\varepsilon)$ necessary and sufficient.
- Recall: the deterministic comm. complexity of computing a mechanism outcome is the log of its menu-size. (BGN'17)

Corollary (G.'18)

For every n there exists $D_n(\varepsilon) = \Theta(\log 1/\varepsilon)$ s.t. for every $\varepsilon > 0$, $D_n(\varepsilon)$ is the minimum communication complexity that satisfies the following: For every distribution $F \in \Delta\big([0,1]^n\big)$ there exists a mechanism M s.t. the deterministic comm. complexity of running M is $D_n(\varepsilon)$ and s.t. $\mathcal{R}ev_M(F) > \mathrm{OPT}(F) - \varepsilon$. (Holds even if F guaranteed to be product of independent distribs.)

Summary: Menu Size as a Function of n and ε

$C(n,\varepsilon)$:			
		poly(n) BILW'14,BGN'17	some ε $(= {}^5\!/_6)$
		poly(n) $\leq$??? $\leq$ exp(n)	arbitrary fixed ε
1 M'81 RZ'83	poly($^1\!/_\varepsilon$) BGN'17,G'18	exp(n) BGN'17	$\varepsilon \to 0$
$n = 1$	arbitrary fixed n	$n \to \infty$	$\varepsilon = {}^1\!/_n \xrightarrow[n\to\infty]{} 0$

Open Question

Is it true that for every n, a menu-size polynomial in n can guarantee 99% of the optimal revenue for any $F \in \Delta(\mathbb{R}_+)^n$?

Multiplicative loss seems the "right goal" for fixed ε and $n \to \infty$.

- For additive ε loss and values in $[0,1]$, the lower bound of Babaioff-G.-Nisan'17 implies that exp(n) menu-size required.
- Somewhat intuitive, as "total welfare in market" may grow linearly with n (and does so in their analysis). Better goal for additive loss is additive $n\varepsilon$ (or equivalently, additive ε when values bounded in $[0, {}^1\!/_n]$) — quite similar to multiplicative ε.

Better core analysis than nudge and round?

Structure in the Core / Improve on Nudge & Round

- "Nudge and round" uses hardly any structural information about the mechanism. Indeed, any mechanism (not necessarily optimal) can be rounded using nudge and round.
- Some results were able to improve upon nudge and round via structural information due to additional assumptions:
  - Dughmi-Han-Nisan'14 achieve near-optimal revenue for monotone valuations (good i is always valued less than good $i+1$, e.g., ad auctions) via polynomial menu-size.
  - Wang-Tang'14: sufficient conditions / families of distributions for which the optimal menu has very small (≤ 4; ≤ 6) size.
  - G.'18 slightly shaves exponent for a standard hazard condition.
  - As we have seen, Babaioff-G.-Nisan'17 achieve separate-selling revenue for independent valuations via polynomial menu-size.
- But structure of optimal mechanism, even for two items, even i.i.d., even bounded, mostly not understood.
- Main open problem: 99% of revenue via poly(n) menu-size, even for i.i.d. items, even for bounded distributions.
- Additional open problem: constructive upper-bound proofs.
  - As opposed to "start with an optimal menu and discretize."

Qualitative Results: Uniform Convergence

Restricting only to limits, most models pretty well understood:

$$\forall n, m: \quad \inf_{F \in \Delta(\mathbb{R}_+^n)} \frac{\mathcal{R}ev_{[m]}(F)}{\mathcal{R}ev(F)} = 0$$

$$\forall n: \quad \inf_{F \in \Delta(\mathbb{R}_+)^n} \frac{\mathcal{R}ev_{[m]}(F)}{\mathcal{R}ev(F)} \xrightarrow[m\to\infty]{} 1$$

$$\forall n, H > L > 0: \quad \inf_{F \in \Delta([L,H]^n)} \frac{\mathcal{R}ev_{[m]}(F)}{\mathcal{R}ev(F)} \xrightarrow[m\to\infty]{} 1$$

$$\forall n: \quad \sup_{F \in \Delta([0,1]^n)} \big(\mathcal{R}ev(F) - \mathcal{R}ev_{[m]}(F)\big) \xrightarrow[m\to\infty]{} 0$$

- As noted, not understood well enough: How fast must C grow as a function of n and ε? (What is the rate of (uniform) convergence?)
- May also be interesting: other restrictions/relaxations that provide uniform convergence/uniform approximation?

More than One Buyer

- Everything so far was for one buyer.
- Completely open: extend above results to $\geq$ one buyer.
- But how should menu-size be defined? The menu faced by each buyer depends on the valuations of the other buyers.
- Largest menu every shown? Sum of menu sizes? Size of union of menus? Something else?
- Or maybe focus on continuing to capture the communication complexity of running the mechanism?
- Related (welfare maximization literature): Dobzinski'16: for rich-enough valuations (far beyond additive):
    - communication complexity $\approx$ log of number of possible menus shown to a buyer ("taxation complexity"),
    - query complexity $\approx$ largest menu shown to any buyer.
    - Does not apply to additive valuations (or even to gross-substitute valuations).
- What does capture these complexities for additive buyers?
- Required complexities for good revenue guarantees?

Summary of Open Questions (that interest me)

- 99% revenue via poly(n) menu-size in any model (even i.i.d, even bounded)?
- Constructive upper bounds? Efficient construction?
- Significantly tighter polynomials? (e.g., $m^{1/7}$ in HN'13)
- Generalizations of above results for multiple buyers?
- Other restrictions (e.g., independent/bounded) or relaxations (e.g., additive) that yield uniform approximation?
- Other auction complexity measures?
- The not-yet-stated question underlying your EC'19 paper!

An Aside: Query Complexity for Complex Valuations

- Dobzinski'11, Dughmi-Vondrak'11, Dobzinski-Vondrak'12, Nisan'14 (survey, proof credited to Dobzinski) study the query complexity of **welfare-approximating** auctions for combinatorial (general, not necessarily additive) valuations.
- General proof scheme:
  1. For any mechanism that guarantees good welfare, there exist a buyer i and valuations for all other buyers s.t. buyer i faces a large menu when all others have these valuations.
  2. Number of value queries to buyer i is $\geq$ menu size she faces.
- Sketch of second step, simplified for deterministic mechanisms and general combinatorial valuations:
  - Fix the valuations of all other buyers. Let buyer i value **bundle** B by the price of bundle B according to the menu.
  - Buyer i is completely indifferent between any two bundles.
  - Consider the scenario where the value of buyer i for a certain bundle $\hat{B}$ was actually one more than the price of that bundle. For the mechanism to rule this out, it must query the value of buyer i for each offered bundle.
- Above papers: restricted valuations, randomized auctions.

Zooming into the Action: a Brief History of Menu Sizes

Multiplicative revenue approximation:

1 item many items (combinatorial valuations)

1 item many items w/ additive valuations many items (combin

1 item 2 items w/ additive valuations many items w/ additive

1 item 2 items w/ independent additive valuations 2 items w/

Precise revenue maximization:

1 item 2 items w/ i.i.d. additive bounded valuations

Kira's lecture after the break!

Questions?

Thank you!

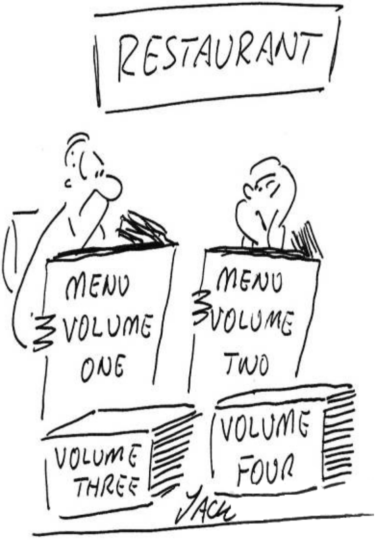

"Lots of choice, isn't there!"